Week 10 · Act 1 — SL shines

Goal: Measure supervised learning on a static, labelled holdout (NSL-KDD via the library).

Thesis: When the world is static and labelled, supervised learning is the right tool.

Run top to bottom. Note the holdout metrics for the exit ticket.

In [ ]:
```
import seas8414_teach as teach
teach.setup()
print(teach.__version__)
teach.story_title_card(
    "Act 1 · Supervised learning shines",
    "Holdout from the SAME distribution as training.",
    ["Labels exist", "One-shot fit", "Stationarity assumed"],
)
```

In [ ]:
```
drift = teach.measure_drift()
print(drift.in_distribution_recall, drift.in_distribution_precision, drift.i
```

In [ ]:
```
teach.story_act1(drift)
```

In [ ]:
```
# Optional: needs OPENAI_API_KEY in ~/.env — skip if unset
teach.ask(
    "Why is near-perfect holdout performance the high-water mark for SL — nc
    recall=drift.in_distribution_recall,
    precision=drift.in_distribution_precision,
    fpr=drift.in_distribution_fpr,
)
```

Exit ticket

What single assumption makes SL the right tool here?